

Literal vs. figurative language: Different or equal?

Rachel Giora

Department of Linguistics, Tel Aviv university, Tel Aviv, 69978, Israel

Abstract

Are literal and nonliteral utterances processed differently or do they follow the same comprehension routes? Relying on intuition, we might expect them to differ. Recent findings, however, do not corroborate this intuition. Evidence from research into moment by moment comprehension demonstrates that such questions are irrelevant to early comprehension processes. And although later integration processes seem more adept at distinguishing literal language from nonliteral language, this conclusion is also unwarranted. Instead, the factor that best accounts for differences occurring early in comprehension is the degree of salience of the instances involved (Giora, Rachel, 1997. Understanding figurative and literal language: The graded salience hypothesis. *Cognitive Linguistics* 7: 183–206; Giora, Rachel, 1999. On the priority of salient meanings: Studies of literal and figurative language. *Journal of Pragmatics* 31, 919–929; Giora, Rachel, in press. *On our mind: Salience, context, and figurative language*. New York: Oxford University Press). Later processes, however, are governed by a functional principle that also does not distinguish between literal nonliteral language. © 2002 Elsevier Science B.V. All rights reserved.

Keywords: Salience; Metaphor; Irony; Literal/nonliteral language; Activation; Retention; Suppression

1. Introduction

The growing interest in context effects, particularly since the inception of pragmatics, has ignited an old debate regarding literal and nonliteral language (for a review of early literature see Reinhart, 1976; for reviews of more recent studies see Ariel, 2002, this issue; Dascal, 1987, 1989; Gibbs, 1994; Giora, 1997, in press). Do literal and nonliteral utterances follow the same processing route or are they processed differently? The question, one should recall, centers on the early ‘moments’ of comprehension. It concerns the factors affecting comprehension immediately upon

E-mail address: giorar@post.tau.ac.il

encountering an utterance. Is it pragmatics (e.g., contextual or world knowledge) or linguistic meaning that determines initial processes? To illustrate the question, let us consider a few examples:

(1) The fisherman needed some fresh fish so he decided to try the west bank.

Is it contextual information—our knowledge of fishing and fish—that plays a crucial role in initial processes, or is it the saliency of the individual words and expressions—e.g., the prominence of the ‘occupied territories’ sense of (the homophone) *West Bank* — that dominates initial comprehension? If it is contextual information that is predominant, then addressees would access the contextually compatible (though less frequent, less-salient) ‘riverside’ meaning of *west bank* exclusively. If it is the accessibility of the words’ meanings that dominates initial processes, then the ‘territorial West Bank’ meaning would be activated initially (in the minds of e. g., Israelis), irrespective of contextual information.

Or take an instance of nonliteral language (2 below). Will a conventional irony such as *Very funny* be processed differently initially in the various contexts (2a and b)? Or will it be processed similarly in both types of context, activating both the ironic and the literal readings, regardless of contextual bias, because it is meaning rather than context that reigns supreme?

- (2) a. Tal and Ortal, the twins, wanted to go to the movies. Their mother recommended a movie she had seen shortly before. When they came home, she was eager to know how they found the movie. They both agreed: “Very funny”.
- b. Iris was walking on her own in the dark alley, when all of a sudden a hand was laid on her back. Startled, she turned around to find out that the hand was her young brother’s who sneaked behind her to frighten her. She said to him: “Very funny”.

The following is a conversational example. In (3), contextual information is biased in favor of the literal meaning of the string of words also constituting an idiom. However, it does not block the contextually incompatible (idiomatic) meaning of the utterance (*totci lo et ha-mic*, a Hebrew idiom, meaning ‘drive him crazy’):

- (3) Iddo and Omri (aged 7; 8, native speakers of Hebrew) are eating supper together. Iddo fetches himself a glass of juice out of the refrigerator.
- Omri: I want to drink too.
- Iddo’s mother: Iddo, *totci lo et ha-mic* (‘take the juice out [of the refrigerator] for him’).
- Iddo (laughingly): ha... ha... *le-hotci lo et ha-mic* (‘take/squeeze the juice out of him’—meaning ‘drive him crazy’).

Example (1) is an instance of literal ambiguity; (2) exemplifies literal vs. ironic uses; (3) illustrates a literal use of an idiom. How are they processed? What is the role played by word and expression meaning? When does context/pragmatics come into play?

2. Different or equal: Initial processes

Two different approaches as to when and how we make sense of utterances have evolved since the 1970s. The different views diverge with respect to the role two

major factors play in comprehension—context vs. linguistic meaning (see also Glucksberg, 2002, this issue).

2.1. *The standard pragmatic model*

Philosophers (e.g., Grice, 1975; Searle, 1979) posit the temporal priority of linguistic input, assuming for the most part that lexical processes are impenetrable and do not interact with contextual information initially (cf. Fodor, 1983). On this view, dubbed ‘The standard pragmatic model’, context comes into play following lexical processes, thus allowing for contextually incompatible meanings to be involved at the initial access phase. Specifically, the standard pragmatic model attributes temporal priority to literal meanings. Accordingly, initial comprehension of literal and nonliteral language should be identical: it should involve accessing the literal meaning of the utterance first. While this would suffice for literal interpretation, nonliteral language would require further processing. For nonliteral language, the initial, obligatory, literal stage would be supplemented by an optional stage of adjustment to contextual information, resulting in the derivation of the contextually compatible meaning.

For example, in the literally biasing context (2a), *Very funny* would be interpreted literally initially. Given contextual fit, no more processes would be initiated. Similarly, in the ironically biasing context (2b), *Very funny* would also be interpreted literally first. Here, however, context misfit would invite further processing, so that *Very funny* would be reinterpreted and adjusted to contextual information. According to the standard pragmatic model, then, the very initial processing phase is identical for both literal and nonliteral language: it involves a literal interpretation. However, unlike literal language, nonliteral language requires extra (and more complex) inferential processes invited by contextual misfit.

2.2. *The direct access view*

The standard pragmatic model has been challenged by cognitive psychologists, psycholinguists, and linguists who assign a primary role in language comprehension to contextual information (e.g., Gibbs, 1994, 2002, this issue; Sperber and Wilson, 1986/95; Vu et al., 1998).¹ This approach, dubbed ‘The direct access view’, assumes that contextual information interacts with lexical processes very early on. Consequently, a rich and supportive context would affect comprehension to a significant extent, so that only intended, contextually compatible meanings would be tapped directly, without involving any contextually incompatible interpretive phase at all.

¹ For a review of a recent version of the direct access view dubbed ‘constraints satisfaction model’, which is specific about the components that make up a ‘strong context’, see McRae et al. (1998). This model, like other direct access models, argues against the encapsulation hypothesis (Fodor, 1983) and have mainly dealt with syntactic ambiguity. The graded salience hypothesis differs from the various versions of the direct access view in that it assumes that lexical accessing is modular and insensitive to processes occurring outside the input system (see also Fodor, 1983).

Thus, in a literal context (2b), *Very funny* would be interpreted only literally (meaning ‘amusing’); in an ironic context (2a), *Very funny* would be interpreted only ironically (meaning ‘annoying’). In similarly strong contexts, then, literal and non-literal language would involve equivalent initial processes that tap contextually appropriate meanings exclusively.²

2.3. *The graded salience hypothesis*

Findings, however, have not been entirely consistent with either view. For instance, contrary to the standard pragmatic model, some instances of nonliteral language (e.g., conventional idioms, conventional ironies, familiar proverbs) have been found to be processed nonliterally initially in nonliterally biasing contexts (e.g., Gibbs, 1980; Turner and Katz, 1997; Giora and Fein, 1999a; Katz and Ferretti, 2001). Similarly, inconsistent with the direct access view, some instances of literal language have been found to be processed nonliterally initially in literally biasing contexts (e.g., Gibbs, 1980; Giora and Fein, 1999a).³

To account for the conflicting findings, I have proposed (Giora 1997, 1999, in press; and see also Turner and Katz, 1997) a more general view of language comprehension—‘The graded salience hypothesis’—which posits the priority of salient (coded, context-independent, prominent) meanings. According to the graded salience hypothesis, salient meanings are processed initially, regardless of either literalness or contextual fit. Context may affect comprehension immediately, but it is ineffective in blocking (contextually incompatible) salient meanings, since it does not interact with lexical processes but runs in parallel (cf. Peleg et al., 2001).

To be *salient*, meanings of words, phrases, or sentences (e.g., the conventional interpretations of idioms or proverbs) have to be coded in the mental lexicon and, in addition, enjoy prominence due to their conventionality, frequency, familiarity, or prototypicality. Meanings not coded in the mental lexicon (e.g., conversational

² It should be noted that the direct access view proposed by Gibbs (1994, 2002, this issue) rejects a strong version of the literal-first model assumed by Grice (1975) according to which the literal interpretation of the utterance as a whole is computed before the nonliteral interpretation is attempted.

³ Indeed, Gibbs’ (1980, 1994, this issue) direct access model rejects only the ‘literal first’ postulate of the standard pragmatic model and does not count as problematic evidence regarding processing of incompatible nonliteral meanings in literally biasing contexts. However, a consistent view of strength of context as primarily affecting comprehension should have identical predictions regardless of contextual bias. Literal context should induce (contextually compatible) literal interpretation exclusively inasmuch as nonliteral context should induce (contextually compatible) nonliteral interpretation exclusively. Findings showing that inappropriate salient (idiomatic) meanings were involved in processing of idioms embedded in literally biasing contexts (Gibbs, 1980) or findings showing that the salient conventional meanings of indirect requests are involved in computing the literal (question) interpretation of these requests (Gibbs, 1983, 1984) are incompatible with the view that attributes a primary role to strength of contextual information in language comprehension. On the contrary, they suggest that contextual information was ineffective in blocking salient but contextually incompatible meanings. Only a direct access view that postulates a salient-first model (Giora, 1997, in press), that is, a direct access of salient meanings (see also Keysar et al., 2000) can account for findings in Gibbs (1980, 1983) and Giora and Fein (1999a).

implicatures constructed on the fly) are *nonsalient*. Coded meanings that are less familiar or less frequent are *less-salient*. Thus, for people living in urban societies, the ‘institutional’ meaning of *bank* is salient and its ‘riverside’ meaning is less-salient; for Internet freaks, the nonliteral meaning of *surf* is salient and its literal meaning may be less-salient. Similarly, for Israelis and Palestinians, the ‘territorial’ sense of *The West Bank* (1 above) would be more salient than for Canadians. A less familiar idiom (*The goose hangs high*) may have a more accessible literal meaning (made up of the salient meanings of its constituents)⁴ and a less accessible idiomatic meaning. Along the same lines, a familiar, conventional irony (*Very funny*; *Big deal*) may have two similarly salient meanings—the literal and the ironic. An innovative use (*Read my lipstick*) of a conventional (fixed) expression (*Read my lips*) may have several salient interpretations—literal (‘Pay attention to what is said’ related to the literal interpretations of *Read my lips*), ironic (‘Discredit what is said’ related to the ironic interpretations of *Read my lips*) as well as several nonsalient interpretations—literal (related to *Read my lipstick*), and ironic (questioning the credibility of the speaker on account of her ‘femininity’). Salience, then, is a matter of degree, determined primarily by frequency of exposure and experiential familiarity with the meaning in question (see Gernsbacher, 1984 and Giora, in press for an elaborate discussion).

Salient meanings are assumed to be accessed immediately upon encounter of the linguistic stimulus via a direct lookup in the mental lexicon. Less-salient meanings will lag behind. Nonsalient meanings require extra inferential processes, and for the most part strong contextual support. Thus, given similarly strong contexts, literal and nonliteral utterances would involve similar initial processes if they are similarly salient. They would not, if they involve salience difference.

Initially, contextual information has a limited role. It may be predictive and avail compatible meanings on its own accord, but it does not affect lexical access and is, therefore, ineffective in blocking (particularly highly) salient meanings.⁵ Though the addressee may rely on all available sources while engaging in comprehension, contextual information does not play a selective role at the initial access stage. Specifically, the prediction of the graded salience hypothesis is that when strength of context and salience are balanced, context will not inhibit the salient but unintended meaning. Even when a context is so highly constrained that it predicts the appropriate meaning very early, even before the linguistic stimulus is encountered, it will not restrict access of salient but contextually incompatible meanings (see Peleg et al.,

⁴ The compositional meaning of a sentence is not salient, since it has not been lexicalized. However, to the extent that the intended, compositional meaning of a sentence relies on the salient meanings of its components, it would be easier to process than when it is not. Though the sum of the components’ meanings is not listed in the mental lexicon as a discrete unit, when integrated into a sentence, this compositional meaning will be more accessible than the compositional meaning of an equivalent sentence whose intended meaning hinges on the less salient meanings of its lexical entries.

⁵ Contextually incompatible meanings accessed initially on account of their salience may be either maintained or suppressed later on depending on their role in constructing the compatible meaning (see later).

2001). To debunk the graded salience hypothesis, highly salient meanings that are contextually implausible should be shown to be obstructed when tapped.⁶

The graded salience hypothesis, then, assumes that, at the initial access phase, literal and nonliteral utterances should not vary processing-wise; they should avail the salient meaning(s) initially, regardless of contextual information or literality. Only when salient meanings are contextually incompatible are additional processes or a strong context required. The distinction, then, that best predicts processing differences is not the literal/figurative divide, but the salient-nonsalient continuum.

2.4. Findings

A review of the evidence adduced by means of various methodologies and materials may enable us to test the different hypotheses. Findings of differences between literal and nonliteral utterances may support the standard pragmatic model. They may also support the graded salience hypothesis in case a salience difference is involved. They would not, however, support the direct access view, since they suggest that a contextually incompatible stage is involved. In contrast, findings of similarity in interpreting literal and nonliteral language may support the direct access view, arguing for equivalent, contextually compatible processes. They may also support the graded salience hypothesis in case utterances are balanced in terms of salience. However, findings of difference between various tropes diverging in salience can be accounted for only by the graded salience hypothesis; they are irrelevant to the standard pragmatic model and are unaccountable by the direct access view.

Given the limited scope of this paper, this review will be rather selective, and focus only on recent exemplars of studies (for a more comprehensive review, see Giora, 1997, in press).

2.4.1. Findings of difference

2.4.1.1. *Reading times.* Reading times of complete utterances have been criticized for being a crude measure, because they have the potential of masking underlying processes (Janus and Bever, 1985; Gibbs and Gerrig, 1989). However, findings of different (as opposed to equal) reading times of statements may be telling about differences involved. Even if reading times cannot reveal the locus of difference, they are, nonetheless, fitted for tapping processes occurring at the message (compositional) level without conflating it with word level meanings (see Gibbs, 2002, this issue).

Recent research has indeed attested to differences obtaining between various types of language use. Studies have shown that unfamiliar (nonsalient) metaphors took

⁶ For a more detailed discussion of how the graded salience hypothesis differs from current accounts of lexical access see Giora (in press), Peleg et al. (2001).

⁷ To be precise, the compositional meaning of an utterance intended literally is not salient because it is not coded in the mental lexicon (cf. note 4). At times, however, it will be treated here as 'salient', because its interpretation hinges on the salient meanings of the components that make it up. Calling a literal interpretation 'salient' is, then, only a shortcut.

longer to read than their literal (salient) interpretations⁷ (Giora and Fein, 1999a), and longer than (salient) literal items (Pexman et al., 2000). They also took longer to read than familiar (salient) metaphors (Gentner and Wolff, 1997; Giora and Fein 1999a; Pexman et al., 2000). Similarly, less familiar (nonsalient) ironies took longer to read than their (salient) literal interpretations (Giora, in press; Giora et al., 1998) and longer than familiar (salient) ironies (Giora and Fein, 1999b). Salient metaphors (*children are precious gems*) took longer to read when embedded in irony inducing contexts (invoking nonsalient ironic meanings) than when embedded in metaphor inducing contexts (activating salient metaphoric meanings) (Gibbs, 1998: 177–180; Pexman et al., 2000). Idioms (*Didn't spill the beans*) took longer to read when embedded in literally biasing contexts (which invite their less accessible literal or compositional interpretation) than when embedded in idiomatically biasing contexts (invoking their salient, idiomatic meanings). Variant (less-salient) idioms (*Didn't spill a single bean*) took longer to read than their salient, (idiomatic) uses (McGlone et al., 1994; Van de Voort and Vonk, 1995). Conventional (salient) metaphors took less time to read when embedded in a context instantiating novel (nonsalient) metaphors rooted in the same conceptual metaphor than when embedded in a context instantiating similarly rooted, but conventional (salient) metaphors (Keysar et al., 2000). Unfamiliar (nonsalient) proverbs took longer to read than their (salient) literal counterparts (Turner and Katz, 1997; Katz and Ferretti, 2001). Unfamiliar (nonsalient) proverbs took longer to process in a figuratively oriented context than in a literally oriented one. While the former invited nonsalient, figurative interpretations, the latter invited more salient, literal interpretations (Honeck et al., 1998). Taken together, these findings attest that it is salience, rather than either contextual information or literality, that determines initial processing. When the compositional meaning of an utterance is what is facilitated (e.g., the literal meaning), this mostly hinges on the salient meanings of the individual words that make it up, rather than on their literality.

2.4.1.2. Lexical decisions. Participants are engaged in lexical decision tasks when, upon encountering a critical (ambiguous, figurative) word in a (context of a) target sentence, they have to make a decision as to whether a probe is a word or a non-word. Faster response times (e.g., pressing a 'yes' or 'no' key) suggest that the concepts responded to were activated during the comprehension process. For example, upon hearing *bugs* in *The man was not surprised when he found several spiders, roaches and other bugs in the corner of his room*, participants responded faster to related words (e.g., 'ant', 'spy') displayed on a monitor screen than to unrelated words (cf. Swinney, 1979). These findings suggest that the related words were activated during the comprehension process. Lexical decision tasks can tell us about underlying initial processes even when reading times of full sentences do not vary.

Using lexical decision tasks, Giora and Fein (1999b) and Giora et al. (1998) showed that novel (nonsalient) ironies and their literal interpretations activated their salient, literal meaning initially (150 ms after offset of the target sentence).⁸ However, in the ironic context, the compatible, nonsalient ironic interpretation required

⁸ Salience was established independently by pretests.

extra processing time and emerged only later on, 1000–2000 ms after offset of the target sentence, though not at the cost of suppressing the contextually incompatible, literal meaning. Where the compatible interpretations were literal, literal targets did not benefit from the extra processing time allowed. Such findings are consistent with the salient and literal first models, but are inconsistent with the direct access view.⁹

Taken together, findings of difference favor the graded salience hypothesis. They show that less salient and nonsalient instances of language took longer to read than their salient counterparts and that salient meanings were activated initially, regardless of either literality or contextual fit.

2.4.2. *Findings of equivalence*

2.4.2.1. *Reading times.* In Giora and Fein (1999a), familiar (salient) metaphors and their literal (salient) uses were shown to take equally long to read. Similarly, Katz and Ferretti (2001) and Turner and Katz (1997) showed that familiar (salient) proverbs and their literal (salient) uses took equally long to read. Though earlier findings attest that, in a rich context, comprehension of novel (nonsalient) metaphors did not take longer to read than their literal uses (Ortony et al., 1978), these equivalent reading times were shown to mask contextually incompatible online processes (Janus and Bever, 1985). Using an online (eye-tracking) methodology, Janus and Bever measured reading times at the locus of the figurative information rather than at the end of the sentence (which is the locus of wrap up processes that might be effortful). Their findings showed that Ortony et al.'s targets exhibited longer reading times when embedded in metaphor inducing contexts (inviting their nonsalient, metaphoric meaning) than when embedded in literal inducing contexts (inviting their salient, literal meaning).

Findings of equal reading times for equally salient utterances may be accounted for by both the graded salience hypothesis and the direct access view. However, more sensitive measures (e.g., lexical decision tasks) are necessary to tap online processes so as to reveal whether contextually incompatible meanings are involved in processes that seem equivalent.

2.4.2.2. *Lexical decisions.* The time taken to make a lexical decision, (e.g., when participants have to make a decision as to whether a related probe is a word or a nonword) upon encountering the target (ambiguous/figurative) word in a given sentence may reveal which concepts were processed online (and responded to fast) and which were not (and took longer to respond to). Using such measures, Giora

⁹ Gibbs (2002, this issue) argues that findings from lexical decision tasks cannot question the direct access model he proposes, since they only attest to activation of word meanings and do not testify to the activation of the overall utterance interpretation. This, however, is not warranted where the time lapse between the offset of the target sentence and the display of the test word is as long as 1000–2000 ms. Such a long delay suggests that the meaning in question has not been only momentarily activated. Nor has it been suppressed as irrelevant (see Swinney, 1979). Rather, it has been retained and integrated with the contextual information and affected interpretation of the sentence as a whole. In addition, the findings from lexical decisions don't stand in isolation, but are corroborated by findings from reading times (see above).

and Fein (1999a) showed that utterances having both salient literal and ironic interpretations (*Very funny*) were processed similarly when embedded in contexts biasing their reading either toward the literal or toward the ironical interpretation: They were initially accessed both ironically and literally, regardless of contextual information. In Williams (1992), polysemies such as *firm*, whose literal and metaphoric meanings are salient, were shown to be processed similarly in both the metaphor inducing context (*The schoolteacher was criticized for not being firm*) and the literal inducing context (*The couple wanted a bed that was firm*). They were initially accessed both metaphorically and literally, regardless of contextual information.¹⁰ These findings are consistent with the salient first model only. They show that salient meanings are activated initially, regardless of context or literality.

In sum, recent research into literal and nonliteral language demonstrates that both the standard pragmatic model and direct access model have limited explanatory power with regard to initial processes. They cannot account for all the findings available. For instance, contrary to the standard pragmatic model, some instances of nonliteral language are processed nonliterally initially (e.g., Van de Voort and Vonk, 1995; Giora and Fein, 1999b). Similarly, contrary to the direct access view, some instances of literal and nonliteral language involve contextually incompatible interpretations initially (e.g., Giora et al., 1998; Giora and Fein, 1999b; Katz and Ferretti, 2001). However, taken together, all of the findings are accountable by the graded salience hypothesis. They show that where salience disparity is involved, different processes ensue, with salient meanings always being accessed faster; where utterances share salience, equivalent processes are operative. In all, salient meanings have not been shown to be preempted by context.

2.4.3. Counterexamples?

2.4.3.1. *Apt metaphors.* At first blush, research into apt metaphors seems to provide a counterexample to the graded salience hypothesis. It suggests that aptness might compensate for low salience. Metaphor aptness (e.g., *That sauna is an oven; That casino is a drug*) is associated with high similarity/shared features (heat; addiction) obtaining between the source/vehicle (*oven; drug*) and the target/topic domains (*sauna; casino*) of the metaphoric expression (Malgady and Johnson, 1976; Johnson and Malgady, 1979; Tourangeau and Sternberg, 1981, 1982; Marschark et al., 1983; Katz, 1986; Tourangeau and Rips, 1991; Blasko and Connine, 1993; Gentner and Wolff, 1997; Chiappe and Kennedy, 1999). Apt metaphors are fast to interpret (Malgady and Johnson, 1976; Johnson and Malgady, 1979; Marschark et al., 1983; Gentner and Wolff, 1997), probably due to the high number of shared features obtaining between the metaphor components.

Blasko and Connine (1993) showed that metaphors rated low in familiarity (salience) and high in aptness were processed faster than low-familiarity/low-aptness metaphors. Particularly, while both types of metaphors activated their literal meanings

¹⁰ There was significant priming only in the literal condition ('solid') which was unaffected by delay. Priming in the metaphoric condition ('strict') "was numerically present at all three delays (decreasing over delay) but the effect was only significant when data were collapsed over delay (Williams, 1992: 201).

instantly, low familiarity/high aptness metaphors also activated their metaphoric meanings instantly. These findings suggest that aptness overrides salience.

However, a close look at the test words used in Blasko and Connine (1993) suggests that this might be an unwarranted conclusion: More than half of the so-called metaphorically related test words they used could be easily primed by the topic rather than by the metaphoric constituent of the utterances.¹¹ When test words are closely associated with the meanings of both the topic and metaphoric constituents, it is no wonder that they are primed instantly. This, however, is not true of the set of low familiarity/low aptness metaphors used, where a smaller number of test words are also highly related to the topic constituents.¹² Indeed, apt metaphors are difficult to test in this way because of the similarity between their target (topic) and source (vehicle) domains (see also Giora and Fein, 1999a; Chiappe et al., submitted).

11

Low familiarity/high apt metaphor (Test words closely associated with the meanings of both the topic and metaphoric constituents are italicized)	Metaphoric test word
---	----------------------

A dusty and crowded attic is a paradise	memories
The <i>fall</i> of an empire is the sunset	decline
A sea <i>captain</i> is a quarterback	leader
Her boyfriend's <i>look of hate</i> was a laser	piercing
Anger was a blizzard	blinding
The <i>thunderclouds</i> were wild horses	rampage
<i>Purgatory</i> is a lobby	waiting
The rocky mountains were a spine	foundation
Perjury is a boomerang	backfires
<i>Indecision</i> was a whirlpool	confusion
<i>Smog</i> is a shroud	engulfing
<i>Greed</i> is a buzzard	consuming

12

Low familiarity/high apt metaphor (Test words closely associated with the meanings of both the topic and metaphoric constituents are italicized)	Metaphoric test word
---	----------------------

A long distance swimmer is a warrior	stamina
A good professor is an oasis	fulfilling
The stars were snowflakes	unique
Ritual is a prison	restricting
A <i>well trained</i> fighter is a knight	skillful
Her thoughts were a boiling kettle	turmoil
Dictionaries are microscopes	detail
Darkness is a gloved hand	unknown
The bible is cement	unchanging
The stars were signposts	navigate
The <i>automobile</i> is a horse	transportation
Garden <i>weeds</i> are a case of the measles	spreading

But aptness has been shown to facilitate processing of less familiar metaphors when other measures were used. For instance, Gentner and Wolff (1997) asked participants to type interpretations of metaphors displayed on a monitor as soon as they had them well formulated. The time between the appearance of the metaphor and the first keystroke was the dependent measure. Either the topic or the vehicle, or both, or none were primed by a mention of each, prior to the display of the metaphor. For instance, before displaying *A job is a jail*, either ‘A job is a jail’ was displayed, which primed both the topic (*job*) and the vehicle (*jail*), or ‘A— is a jail’ was displayed, which primed only the vehicle, or ‘A job is a—’ was displayed, which primed only the topic, or ‘A—is a—’ was displayed, which primed none. Across the board, apt metaphors were processed faster than less apt ones. However, apt but less familiar metaphors were not processed as fast as familiar metaphors, suggesting that, though aptness is a factor in metaphor comprehension, it does not override salience.^{13,14}

2.4.3.2. Metaphor-based polysemies. Some recent accounts abandon the view that, overall, words activate specified senses initially (compare Fodor’s, 1983: 86 “shallow output”). They propose, instead, the underspecification model, which should be particularly relevant to polysemies—words that have multiple related senses. The meanings of such words, they contend, are stored as highly abstract, underspecified entities whose specific meanings are determined by contextual information (Groefsema, 1995; Ruhl, 1989; Frazier and Rayner, 1990).

Following Frazier and Rayner (1990), Pickering and Frisson (2001) applied the underspecification model to account for how polysemous verbs are processed. Polysemous verbs (e.g., *disarm* meaning ‘reduce weapons’, ‘neutralize’), they argue, should be processed differently from ambiguous verbs—verbs having multiple unrelated meanings (e.g., *dock* meaning ‘boat landing’, ‘decrease’). While ambiguities force a selection between inconsistent meanings, a polysemy need not be disambiguated instantly, because the multiple senses it involves are not mutually incompatible. Upon encountering a polysemous verb, the processor activates a single, underspecified

¹³ Though Gildea and Glucksberg (1983) and Glucksberg et al. (1982) showed that apt metaphors are understood instantly and require no contextual information, whereas less apt metaphors rely on contextual information for their interpretation, they did not look into the salience/familiarity variable.

¹⁴ That aptness might not be a factor when familiar metaphors are concerned might be deduced from Pynte et al. (1996). Pynte et al. tapped processes involved in what they term familiar and unfamiliar metaphors by measuring brainwaves. It is hard, however, to draw conclusions from their study as to how salience (or familiarity) affects processing since, in fact, they only used familiar metaphor vehicles (e.g., *lions*). Consequently, their ‘familiar’ (*Those fighters are lions*) and ‘unfamiliar’ (*Those apprentices are lions*) metaphoric utterances differ only in aptness, that is, in how similar/relevant their vehicle is to their topic. Out of context, ‘familiar vs. ‘unfamiliar’ items did not differ (Experiment 2). However, when the targets (*lions*) were incompatible with prior context (“They are not idiotic: Those fighters are lions”) they elicited larger N400 amplitude than when cohering with context (“They are not cowardly: Those apprentices are lions”) (Experiments 3 and 4). At best, these findings can be taken to suggest that (relatively) familiar metaphors involve their salient metaphoric meaning upon encounter, regardless of either context or aptness.

meaning first. When resolution is required, it uses contextual information to home in on the precise, contextually appropriate sense. Results from eye-tracking measures indeed show that initial processing of polysemous verbs did not demonstrate any early effects of either frequency or context (which, however, emerged downstream).

The underspecification account may be inconsistent with the graded salience hypothesis, which seems to assume that initial access involves discrete senses. It is also inconsistent with findings such as Williams' (1992) and Garrod and Terras' (2000), which demonstrate that polysemies activate specific senses initially. However, lack of early frequency effects requires some explanation. Though frequency is a component of salience, it is not the only factor determining salience. Given that the metaphors selected by Frisson and Pickering are highly conventional and familiar, it is quite possible that both senses of the polysemous verbs (the literal and the metaphoric) are similarly salient. Consequently, both would be activated initially.

Alternatively, given the fact that the locus of these verbs' resolution is further downstream, it is plausible that the processor activates an underspecified meaning initially and, in the course of time, zooms in on the salient sense(s) first, with the less-salient sense(s) reaching sufficient levels of activation somewhat later. Lack of early effects might indeed become pronounced when there is no need to select between (salient and less-salient) incompatible senses and suppress the contextually inappropriate sense. In any event, the underspecification model is consistent with the assumption of the graded salience hypothesis that (similarly salient) literal and figurative utterances should be initially processed similarly, regardless of either context or literality.

2.4.3.3. Mapping complexity. In a recent study, Coulson and Van Petten (submitted) proposed that mapping complexity plays a crucial role in language comprehension. Accordingly, the literality/metaphoricity dichotomy should give way to a complexity continuum. To test their view, Coulson and Van Petten used N400 amplitude measures. N400 is a brainwave, which is large at the beginning of a sentence, particularly for low-frequency words, but declines with increasing semantic constraints as the sentence proceeds (Van Petten, 1995). Among other things, N400 amplitude measures are sensitive to ease or difficulty of meaning construction. Using such measures, Coulson and Van Petten show that not all literal meanings are alike. Some are less 'literal' than others, involving more complex processes (termed literal mappings) typically attributable to metaphors. In their study, sentence-final words (*gem*) were matched across conditions for cloze probability, word length, and word frequency, diverging, however, along a complexity continuum. The literal sense of the target word (*That stone we saw in the natural history museum is a gem*) involved simple processes, and conveyed its conventional, literal meaning. The metaphoric sense (*After giving it some thought, I realized the new idea is a gem*) involved most complex processes linking the speaker's idea analogically to a gemstone to evoke its brightness and clarity. The intermediate case (*The ring was made of tin, with a pebble instead of a gem*) involved literal mappings. It prompted readers to map a conceptual structure from a different domain. Processing *gem* in its conventional sense under this condition exploited a correspondence between a worthless toy ring and

the more prototypical expensive ring—both construed as rings, both worn on the finger, and both having a small, roundish, hard object in the center. Results indeed support the hypothesized continuum. They show that while metaphors elicited larger N400s than literal sentences, literal mappings fell between metaphors and literals.

These findings are indeed consistent with a continuum hypothesis that suggests that processing complexity is a crucial factor in comprehension. However, a close look at the items used may allow for an alternative, not mutually exclusive explanation. A review of the items reveals that quite a few of those embedded in the ‘literal mapping’ sentences were used in their less-salient/less prototypical (though conventional, i.e., coded) meanings. For instance, in the intermediate condition exemplified above (*The ring was made of tin, with a pebble instead of a gem*), it is the less-salient (‘small, roundish form’) sense of *gem* that is intended rather than its salient ‘brightness’ ‘valuable’ or ‘decorative’ meaning invited by the literal sentence-context (see also Ortony et al., 1985). In addition, many of the metaphoric items do not have salient, metaphoric meanings either, explaining their difficulty of processing. It is thus possible that in addition to mapping complexity, the items also diverge on meaning salience, which may explain the differences found.

3. Different or equal: Integration processes

Data regarding early processes reveal that it is salience rather than either context or literality or nonliterality that determines the processes involved initially. If the salient meaning accessed initially is contextually compatible, no more activation processes are invited. If it does not meet contextual fit, additional activation processes are required. When literal and nonliteral utterances converge in degree of salience, they involve the same processes initially (e.g., familiar metaphors and their literal counterparts; familiar proverbs and their literal counterparts). When the meanings of literal and nonliteral utterances diverge in salience, they are accessed differently (e.g., unfamiliar metaphors and their literal interpretations; unfamiliar proverbs and their literal interpretations): Though they both access their salient meanings first, the unfamiliar (nonsalient) utterances require additional activation processes. It is degree of salience, then, that makes a difference.

If early processes do not distinguish literal from nonliteral language, would later integration processes disclose a difference? This question is particularly intriguing regarding similarly salient utterances. Would utterances converging in salience (involving similar initial processes) also exhibit similar processes at the later integration phase? The question is highly relevant with respect to salient but contextually incompatible meanings.

Recall that the assumption that salient meanings are activated initially, regardless of contextual information, implies that inappropriate meanings would be involved in the process just because of their saliency. Thus, processing a conventional irony such as *Very funny* in the ironically biasing context (2a) would involve activating the salient, inappropriate, literal sense (‘amusing’) alongside its salient, appropriate, ironic sense (‘annoying’). By the same token, the salient, inappropriate, ironic sense

would be activated in the literally biasing context (2b) alongside its salient, appropriate, literal sense. Would such inappropriate meanings be discarded as ‘irrelevant’, as might be deduced from studies showing suppression of contextually incompatible meanings (cf. Swinney, 1979; Gernsbacher, 1990; Gernsbacher and Robertson, 1999; Glucksberg et al., 2001; Gernsbacher et al., in press)?

According to the retention hypothesis supplementing the graded salience hypothesis (Giora and Fein, 1999a; Giora, in press), contextually incompatible meanings would not be discarded if they are not disruptive or have a role in constructing the compatible meaning. Thus, the salient ‘amusing’ sense induced by the conventional irony *Very funny* would be retained in the ironically biasing context (2b), because it is conducive to the interpretation of the compatible ironic sense (cf. Giora, 1995). In contrast, the ironic ‘annoying’ sense should be discarded in the literally biasing context (2a), because, in this context, it interferes with the compatible meaning. Similarly, conventional metaphors (*firm*) which are activated initially both literally (‘solid’) and figuratively (‘strict’) should retain the literal sense in the metaphorically biased context, but suppress the metaphoric sense in the literally biased context. On the retention hypothesis, then, later processes should distinguish between literal and figurative utterances converging in salience. While figurative language retains so-called inappropriate (literal) meanings, some instances of literal language involve suppressing incompatible (figurative) senses.

It is quite possible, however, that some instances of polysemies and apt metaphors would be processed similarly even at the later integration phase. Given that their literal and figurative meanings are highly similar, both meanings might be appropriate in both contexts. For instance, while the literal (‘reduce weapons’) sense of *disarmed* should be retained in the figurative context because of its functionality, the figurative (‘weaken’ ‘neutralize’) sense need not be suppressed in the literal context either, since it is compatible with that sense as well.

The direct access view does not predict any differences. Given that a strong context governs interpretation, in a rich ecology, only contextually compatible meanings would be activated and retained for further processes. Integration processes, then, would involve only appropriate meanings, with no recourse to suppression. In contrast, according to the standard pragmatic model, the inappropriate literal meaning of a figurative utterance should be rejected and replaced by the appropriate meaning, hence suppression is expected only in the figuratively biased context. No such operation is anticipated in the literally biasing context, since in that context, the figurative meaning is not expected to emerge. According to the standard pragmatic model, then, literal and nonliteral language would diverge at the integration phase as well.

3.1. Findings of difference following similar initial processes

3.1.1. Lexical decisions

Comparing lexical decisions made immediately and after a long delay may reveal which meanings activated initially are retained in which contexts. Indeed, as predicted by the retention hypothesis (Giora and Fein, 1999a; Giora, in press), inappropriate meanings conducive to the compatible interpretation were found to be

retained even after a long delay. In contrast, meanings conflicting with the compatible meaning were discarded. For instance, in a figurative context (*firm teacher*), where the incompatible literal ('solid') sense is supportive of the compatible meaning ('strict'), it remains active even when the contextually appropriate sense has been determined. In contrast, in the literally biasing context (*firm bed*), where the incompatible metaphoric ('strict') meaning is irrelevant to the compatible meaning ('solid'), it has been deactivated (Williams, 1992).

3.1.2. *Word fragment completion*

In word fragment completion tests, participants are asked to complete a fragmented word (e.g., b-tt-r) with the first word that comes to mind (e.g., *better*, *bitter*, *butter*). Being an offline measure, word fragment completion may disclose the meanings that are retained following the initial access phase. As predicted by the retention hypothesis, findings show that the contextually incompatible, literal meaning of a familiar metaphor is retained in the metaphoric context, since in that context it is conducive to the compatible meaning. In contrast, in the literally biasing context, the contextually incompatible metaphoric meaning is less active, since in that context it interferes with the compatible meaning. (Giora and Fein, 1999a).

3.1.3. *Cued recall*

A cued recall procedure is an offline task that may tap later integration processes. In this memory task, a cue (a test word) is employed that is related to a particular meaning of the target. The assumption is that such a cue would be an effective recall aid in case the meaning it is associated with has been generated during comprehension (cf. Tulving, 1983). Turner and Katz (1997) used cued recall to find out which meanings were processed in which contexts. The stimuli were familiar and unfamiliar proverbs. For instance, for a literal interpretation of the familiar proverb *The grass is always greener on the other side of the fence*, the cue was 'pasture'; for the figurative interpretation of that proverb, the cue was 'envy'. Results from the cued recall test (which followed a reading session of proverbs in literally and figuratively biasing contexts) revealed that familiar proverbs and their literal counterparts (which had taken equally long to read) exhibited somewhat different patterns in the offline task. The literal meaning of the proverb embedded in a figuratively biasing context was recalled to a greater extent than the figurative meaning of the proverb embedded in a literally biasing context. Interpreted in conjunction with the equal time taken to read the proverbs in both contexts, these recall patterns may be viewed as suggesting that the figurative meaning of the proverb was activated initially but underwent deactivation in the context where it was irrelevant to the compatible meaning (this interpretation of the findings is not proffered by the authors). Findings from measures tapping later processes, then, suggest that literal and nonliteral language diverge at the integration phase.

3.1.4. *Counterexample*

Finding of late equivalent processes of literal and nonliteral targets following equivalent initial processes may provide a counterexample to the retention hypothesis.

They show that salient but contextually incompatible meanings that are not involved in constructing the compatible meaning may nevertheless be retained. It is plausible, though, that some meanings are so salient that they are less amenable to suppression, even when they are not instrumental and might divert attention (cf. example 2 above).

For example, Giora and Fein (1999b) found that utterances such as *Very funny*, whose literal and ironic meanings are salient, activated both their literal and ironic meanings initially (150 ms after offset of the target sentence) and also later on (1000 ms after offset of the target sentence) in both types of context. Such findings, while attesting to equal initial processes of similarly salient meanings, nevertheless demonstrate that, contra the retention hypothesis, the salient but inappropriate ironic sense was not deactivated in the literally biasing context.

3.2. Findings of difference following different initial processes

Discussing findings of difference following different initial processes would not shed further light on the question of whether literal and nonliteral language involve different or equal processes. Apparently they do. If initial processes are different, then the utterances in question are processed differently. There is, however, one comparison that might still be worth looking at. Comparing similarly salient meanings in different types of context might be revealing with respect to differences following early processes that are not predicted by these self-same processes. On the retention hypothesis, similarly salient meanings that are accessed initially, regardless of context, should be either retained or suppressed, depending on the part they play in constructing the intended meaning of the utterance. It would therefore be interesting to see whether similarly salient meanings behave differently vis à vis their functions, regardless of contextual appropriateness. Would the literal meanings of both less familiar (*What a lovely day for a picnic* said on a stormy day) and familiar (*tell me about it*) ironies, which are similarly salient in both the ironic and the literal contexts (though not equally compatible), behave differently in both types of context, as would be predicted by the retention hypothesis? Would the salient, literal meanings of novel metaphors and the salient, nonliteral meanings of familiar metaphors behave differently following early processes?

3.2.1. Lexical decisions

Findings in Giora and Fein (1999b) show that salient, though contextually incompatible, meanings that are instrumental in constructing the utterance interpretation are retained longer than equivalent and contextually compatible meanings that are no longer instrumental. For example, the salient, literal meaning ('nice') of (less familiar) ironies (*what a lovely day for a picnic*) is retained longer in an irony-inducing context (a stormy day) than in a literally biased context (a sunny day), in which it is contextually compatible. While the literal meaning may be involved in constructing the ironic interpretation even 2000 ms after offset of the target sentence, this is not the case when this meaning is compatible with (the literal) context. The literal meaning of an utterance embedded in a literally biased context, it is

hypothesized, is comprehended instantly. Having been accessed and integrated, the salient, contextually compatible meaning need not maintain its initial activation level and may begin to fade. Such findings are consistent with the retention hypothesis.

4. Conclusions

Reviewing recent studies on initial processes of literal and nonliteral language, it seems safe to conclude that the factor that best explains the plethora of findings is the degree of meaning salience involved. The bulk of the evidence demonstrates that, regardless of either literality or contextual information, utterances converging in meaning salience exhibit equivalent processes; utterances diverging in meaning salience exhibit different processes. The question of whether literal and nonliteral language are equal or different is, therefore, irrelevant to initial processes.

Findings tapping processes following the initial access stage seem, however, to point to a different conclusion. They suggest that later processes might be governed by a functional principle that might distinguish literal from nonliteral language. They show that while nonliteral language retains contextually incompatible meanings due to their role in constructing the intended meaning (e.g., literal meanings in nonliteral contexts), literal language does not.

This, however, is a premature conclusion. Though no research has looked into the poetics of literal language (but see Kronrod, et al., 2000; Giora, in press), there is no reason to assume that the functional hypothesis should not apply to literal language as well. It is quite plausible that salient, contextually incompatible meanings that are nevertheless conducive to the interpretation of an utterance should be retained regardless of (non)literalness. Consider the example of a shoe shop in a mall named *Body and sole*. Apparently, the contextually incompatible ‘soul’ sense of the literal ambiguity is accessed due to its salience and is not deactivated despite its contextual incompatibility. Given that it is conducive to the construction of the compatible (literal) meaning, it is probably retained for further processes. Similar examples abound. Consider *Love at first site*—an advertisement for an Internet service—which, albeit literal, activates and retains ‘Love at first sight’ alongside *Love at first site*. This must also be true of *Lovett first sight*—the title of an (old) article discussing Julia Roberts’ and Lyle Lovett’s love and wedding. Or, take *Don’t leave without a good buy* inscribed on an airport shop which relates *a good buy* to ‘a good bye’; or *Sofa so good*—a name of a furniture store, which invokes ‘So far so good’ and retains it for comparison purposes. Such examples suggest that activating salient but incompatible meanings that are retained (rather than suppressed) need not distinguish literal from nonliteral language.

Complementarily, it is quite possible that literal ambiguities involving salient meanings that interfere with comprehension would be subjected to suppression. For example, the ‘spy’ sense related to *bug*, which is activated initially due to its salience even in an inappropriate context (*The man was not surprised when he found several spiders, roaches and other bugs in the corner of his room*), is suppressed shortly

afterwards, since it is not conducive to the compatible interpretation (cf. Gernsbacher, 1990; Swinney, 1979, and others). Our intuition, then, that literal and non-literal language involve different processes has not yet gained empirical support. Saliency and functionality, however, seem better predictors of the differences and similarities found.

Acknowledgements

I wish to thank Mira Ariel, Seana Coulson, and Albert Katz for their very helpful comments.

References

- Ariel, Mira, 2002. The demise of a unique literal meaning. *Journal of Pragmatics* 34(4): 361–402.
- Ariel, Mira, submitted. Privileged interactional interpretations. *Journal of Pragmatics*.
- Blasko, G. Dawn and Cynthia Connine, 1993. Effects of familiarity and aptness on metaphor processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 19: 295–308.
- Chiappe, Dan L. and John M. Kennedy, 1999. Aptness predicts preference for metaphors or similes, as well as recall bias. *Psychonomic Bulletin & Review* 6: 668–676.
- Chiappe, Dan L., John M. Kennedy and Penny Chiappe, submitted. Aptness is more important than comprehensibility in predicting recognition bias and preference for metaphors and similes. *Psychonomic Bulletin & Review*.
- Coulson, Seana and Van Petten, submitted. Conceptual integration and metaphor: An event-related brain potential study. *Memory & Cognition*.
- Dascal, Marcelo, 1987. Defending literal meaning. *Cognitive Science* 11: 259–281.
- Dascal, Marcelo, 1989. On the role of context and literal meaning in understanding. *Cognitive Science* 13: 253–257.
- Fodor, Jerry, 1983. *The modularity of mind*. Cambridge: MIT Press.
- Frazier, Lyn and Keith Rayner, 1990. Taking on semantic commitments: Processing multiple meanings vs. multiple senses. *Journal of Memory and Language* 29: 181–200.
- Garrod, Simon and Melody Terras, 2000. The contribution of lexical and situational knowledge to resolving discourse roles: Bonding and resolution. *Journal of Memory and Language* 42: 526–544.
- Gentner, Dedre and Phillip Wolff, 1997. Alignment in the processing of metaphor. *Journal of Memory and Language* 37: 331–355.
- Gernsbacher, Morton Ann, 1984. Resolving twenty years of inconsistent interactions between lexical familiarity and orthography, concreteness, and polysemy. *Journal of Experimental Psychology: General* 113: 256–281.
- Gernsbacher, Morton Ann, 1990. *Language comprehension as structure building*. Hillsdale, NJ: Erlbaum.
- Gernsbacher, Morton Ann and Rachel R.W. Robertson, 1999. The role of suppression in figurative language comprehension. *Journal of Pragmatics* 31: 1619–1630.
- Gernsbacher, Morton Ann, Boaz Keysar and Rachel W. Robertson, in press. The role of suppression and enhancement in understanding metaphors. *Journal of Memory and Language*.
- Gibbs, Raymond W. Jr., 1980. Spilling the bean on understanding and memory for idioms in conversation. *Memory & Cognition* 8: 449–456.
- Gibbs, Raymond W. Jr., 1983. Do people always process the literal meanings of indirect requests? *Journal of Experimental Psychology: Learning Memory, and Cognition* 9: 24–533.
- Gibbs, Raymond W. Jr., 1984. Literal meaning and psychology theory. *Cognitive Science* 8: 575–604.
- Gibbs, Raymond W. Jr., 1994. *The poetics of mind*. Cambridge: Cambridge University Press.
- Gibbs, Raymond W. Jr., 1998. Counterpoint Commentary. In: A. Katz, C. Cacciari, R. Gibbs and T. Turner, eds., *Figurative language and thought*, 158–192. New York: Oxford University Press.

- Gibbs, Raymond W. Jr., 2002. A new look at literal meaning in understanding what is said and implicated. *Journal of Pragmatics* 34(4): 457–486.
- Gibbs, Raymond W. Jr. and Richard J. Gerrig, 1989. How context makes metaphor comprehension seem 'special'. *Metaphor and Symbolic Activity* 4: 145–158.
- Gildea, Patricia and Sam Glucksberg, 1983. On understanding metaphor: The role of context. *Journal of Verbal Learning and Verbal Behavior* 22: 577–590.
- Giora, Rachel, 1995. On irony and negation. *Discourse Processes* 19: 239–264.
- Giora, Rachel, 1997. Understanding figurative and literal language: The graded salience hypothesis. *Cognitive Linguistics* 7: 183–206.
- Giora, Rachel, 1999. On the priority of salient meanings: Studies of literal and figurative language. *Journal of Pragmatics* 31: 919–929.
- Giora, Rachel, in press. On our mind: Salience, context, and figurative language forthcoming. New York: Oxford University Press.
- Giora, Rachel and Ofer Fein, 1999a. On understanding familiar and less-familiar figurative language. *Journal of Pragmatics* 31: 1601–1618.
- Giora, Rachel and Ofer Fein, 1999b. Irony: Context and salience. *Metaphor and Symbol* 14: 241–257.
- Giora, Rachel, Ofer Fein and Tamir Schwartz, 1998. Irony: Graded salience and indirect negation. *Metaphor and Symbol* 13: 83–101.
- Glucksberg, Sam, 2001. Understanding figurative language: From metaphor to idiom. New York: Oxford University Press.
- Glucksberg, Sam, Patricia Gildea and Howard B. Bookin, 1982. On understanding nonliteral speech: Can people ignore metaphors? *Journal of Verbal Learning and Verbal Behavior* 21: 85–98.
- Glucksberg, Sam, Mary R. Newsome and Yevgeniya Goldvarg, 2001. Inhibition of the Literal: Filtering metaphor-irrelevant information during metaphor comprehension. *Metaphor and Symbol* 16: 277–293.
- Grice, H. Paul, 1975. Logic and conversation. In: P. Cole and J. Morgan, eds., *Speech acts. Syntax and semantics* Vol. 3, 41–58. New York: Academic Press.
- Groefsema, Marjolein, 1995. Can, may, would, and should: A relevance-theoretic account. *Journal of Linguistics* 31: 53–79.
- Honeck, P. Richard, Jeffrey Welge and Jon G. Temple, 1998. The symmetry control in tests of the standard pragmatic model: The case of proverb comprehension. *Metaphor and Symbol* 13: 257–273.
- Janus, R.A. and Tom G. Bever, 1985. Processing of metaphoric language: An investigation of the three stage model of metaphor comprehension. *Journal of Psycholinguistic Research* 14: 473–487.
- Johnson, Michael G. and Robert G. Malgady, 1979. Some cognitive aspects of figurative language: Association and metaphor. *Journal of Psycholinguistic Research* 8: 249–265.
- Katz, N. Albert, 1986. On choosing the vehicles for metaphors: Referential concreteness, semantic distances, and individual differences. *Journal of Memory and Language* 28: 486–499.
- Katz, Albert N. and Todd Ferretti, 2001. Moment-by-moment reading of proverbs in literal and non-literal contexts. *Metaphor and Symbol* 16: 193–221.
- Keysar, Boaz, Yeshayahu Shen, Sam Glucksberg and William S. Horton, 2000. Conventional language: How metaphorical is it?. *Journal of Memory and Language* 43: 576–593.
- Kronrod, Ann, Rachel Giora and Ofer Fein, 2000 (September). Creative writing: The optimal creative innovation in fixed expressions. The EARLI special interest group writing conference 2000. University degli Studi di Verona.
- Malgady, Robert G. and Michael G. Johnson, 1976. Modifiers in metaphors: effects of constituent phrase similarity on the interpretation of figurative sentences. *Journal of Psycholinguistic Research* 5: 43–52.
- Marschark, Marc, Albert N. Katz and Allen Paivio, 1983. Dimensions of metaphor. *Journal of Psycholinguistic Research* 12: 17–39.
- McGlone, Matthew, Sam Glucksberg and Cristina Cacciari, 1994. Semantic productivity and idiom comprehension. *Discourse Processes* 17: 167–190.
- McRae, Ken, Michael J. Spivey-Knowlton and Michael K. Tanenhaus, 1998. Modeling the influence of thematic fit (and other constraints) in on line sentence comprehension. *Journal of Memory and Language* 38: 283–312.
- Ortony, Andrew, Diane L. Schallert, Ralph E. Reynolds and Stephen J. Antos, 1978. Interpreting metaphors

- and idioms: Some effects of context on comprehension. *Journal of Verbal Learning and Verbal Behavior* 17: 465–477.
- Ortony, Andrew, Richard J. Vondruska, Mark A. Foss and Lawrence E. Jones, 1985. Saliency, similes, and asymmetry of similarity. *Journal of Memory and Language* 24: 569–594.
- Peleg, Orna, Rachel Giora and Ofer Fein, 2001. Saliency and context effects: Two are better than one. *Metaphor and Symbol* 16: 173–192.
- Pexman Penny, Todd Ferretti and Albert N. Katz, 2000. Discourse factors that influence irony detection during on-line reading. *Discourse Processes* 29: 201–222.
- Pickering, Martin and Steven Frisson, 2001. Processing of verbs: Evidence from eye movements. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 27: 556–573.
- Pynte, Joel, Mireille Besson, Fice-Henri Robinchon and Jezabel Poli, 1996. The time-course of metaphor comprehension: An event-related potential study. *Brain and Language* 35: 293–316.
- Reinhart, Tanya, 1976. On understanding poetic metaphor. *Poetics* 5: 383–402.
- Ruhl, Charles, 1989. *Monosemy: A study in linguistic semantics*. Albany: State University of New York Press.
- Searle, John, 1979. *Expression and meaning*. Cambridge, England: Cambridge University Press.
- Schwoebel, John, Shelly Dews, Ellen Winner and Kavitha Srinivas, 2000. Obligatory processing of the literal meaning of ironic utterances: Further Evidence. *Metaphor and Symbol* 15: 47–61.
- Sperber, Dan and Deirdre Wilson, 1986/1995. *Relevance: Communication and cognition*. Oxford: Blackwell.
- Swinney, David A., 1979. Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior* 18: 645–659.
- Tourangeau, Robert J. and Lance Rips, 1991. Interpreting and evaluating metaphors. *Journal of Memory and Language* 30: 454–472.
- Tourangeau, Roger and Robert J. Sternberg, 1981. Aptness in metaphor. *Cognitive Psychology* 13: 27–55.
- Tourangeau, Roger and Robert J. Sternberg, 1982. Understanding and appreciating metaphor. *Cognition* 11: 203–244.
- Tulving, Endel, 1983. *Elements of episodic memory*. Oxford: Oxford University Press.
- Turner, Nigel E. and Albert N. Katz, 1997. Evidence for the availability of conventional and of literal meaning during the comprehension of proverbs. *Pragmatics and Cognition* 5: 203–237.
- Van Petten, Cyma, 1995. Words and sentences: Event-related brain potential measures. *Psychophysiology* 32: 511–525.
- Van de Voort, Marlies E.C. and Wietske Vonk, 1995. You don't die immediately when you kick an empty bucket: A processing view on semantic and syntactic characteristics of idioms. In: M. Everaert, E.-J. van der Linden, A. Schenk and R. Schreuder, eds., *Idioms: Structural and psychological perspectives*, 283–299. Hillsdale, NJ: LEA.
- Vu, Hoang, George Kellas and Stephen T. Paul, 1998. Sources of sentence constraint in lexical ambiguity resolution. *Memory & Cognition* 26: 979–1001.
- Williams, John N., 1992. Processing polysemous words in context. Evidence from interrelated meanings. *Journal of Psycholinguistic Research* 21: 193–218.

Rachel Giora is Associate Professor of Linguistics at Tel Aviv University. Her research areas include discourse coherence, relevance, cognitive pragmatics, language and ideology, and women and language. Her recent work focuses on figurative language (particularly, irony, joke, and metaphor) and on the notion of saliency.